

# Qi Lin

Tempe, AZ | Email: [qlin36@asu.edu](mailto:qlin36@asu.edu) | Phone: (949) 591-9768 | [GitHub](#) | [Google Scholar](#)

## RESEARCH INTERESTS

---

- **Data Infrastructure/Storage:** Key-value stores (**RocksDB**, **LevelDB**), **RDMA**-based high-speed networking and remote memory access, and high-performance storage devices (**NVMe**, **SPDK**).
- **RDMA and CXL-Based Disaggregated Memory Systems:** Cache-coherent shared memory over **RDMA** and **CXL**, memory consistency and ordering, remote load/store semantics, and failure-resilient shared-memory system design.

## EDUCATION

---

**Ph.D. Student in Computer Science**, Arizona State University. Advisor: [Prof. Zhichao Cao](#) 2023 – 2027 (Expected)  
**M.S. in Computer Science**, University of California, Irvine. Advisor: [Prof. Faisal Nawab](#) 2021 – 2023  
**M.Eng. in Traffic Engineering**, Southeast University 2017 – 2020  
**B.E. in Traffic Engineering**, Chongqing Jiaotong University 2013 – 2017

## WORK EXPERIENCE

---

**Research Intern (Incoming)**, Microsoft Research, Redmond, US Summer 2026  
**Full-Time Algorithm Engineer**, Hikvision, Hangzhou, China Jul 2020 – Sep 2021

## RECENT PUBLICATIONS

---

- **[SIGMOD'26]** Qi Lin, Gangqi Huang, Te Guo, Chang Guo, Viraj Thakkar, Zichen Zhu, Jianguo Wang, Zhichao Cao. "O3-LSM: Maximizing Disaggregated LSM Write Performance via Three-Layer Offloading." *Proceedings of ACM Conference on Management of Data (SIGMOD)*, 2026.
- **[EMNLP'25]** Yuhang Chen, Qi Lin, Zhichao Cao, Tianlong Chen. "Bit-Flip Error Resilience in LLMs: A Comprehensive Analysis and Defense Framework." *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2025.
- **[EDBT'24]** Lixi Zhou\*, Qi Lin\*, Jia Zou. "Serving Deep Learning Model in Relational Databases." *Proceedings of the International Conference on Extending Database Technology (EDBT)*, 2024.
- **[TKDE'24]** Qi Lin, Binbin Gu, Faisal Nawab. "RollStore: Hybrid Data Indexing for Decentralized Blockchain Applications." *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 2024.

## ONGOING RESEARCH

---

- **[HPDC'26]** Qi Lin, Zhenyu Zhang, Viraj Thakkar, Mei Zheng, Zhichao Cao. "An LLM Agent-Driven Automatic Tuning Framework for Data Systems." *Proceedings of the ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC)*, 2026.
- **[Under Preparation for SIGMOD'27]** Zhang Cao, Qi Lin, Chang Guo, Zhichao Cao. "Partial Failure Resilience for Data Systems over CXL-based Disaggregated Memory."
- **[SIGMETRICS'26]** Viraj Thakkar, Qi Lin, Zhichao Cao. "LLM-Automated Optimization Framework over Key-Value Stores." *Proceedings of ACM Conference on Management of Data (SIGMOD)*, 2026.

## RESEARCH EXPERIENCE

---

### **CANON: Partial Failure Resilient Shared Memory System over CXL-Based Disaggregated Memory**

- Designed and implemented a high-concurrency shared-memory system over **CXL-based disaggregated memory**. Built a cache-coherent remote memory architecture where multiple compute threads directly access a shared memory pool via load/store semantics and atomic primitives (**CAS**, **FAA**).
- Addressed **CXL-specific memory ordering challenges** arising from CPU cache hierarchy interaction. Enforced correct visibility of remote updates using **mfence** and **clwb**, and designed atomic update protocols resilient to write reordering and partial writes.
- Formalized structural and semantic invariants for concurrent shared-memory data structures and developed lightweight recovery mechanisms for **partial client failures**, avoiding heavy transaction-based logging overhead.
- Evaluated under multi-threaded workloads on NUMA-based CXL simulation and RDMA cluster environments, demonstrating stable behavior under injected failures and up to 50% throughput improvement over baselines.

### **Exploring Disaggregated Memory for Optimizing LSM-based Key-Value Stores**

- Designed an LSM-based KV store that leverages **Disaggregated Memory (DM)** via **RDMA** to scale writes, achieved **4.5×** higher write throughput and **76%** lower P99 latency.

- Built a **DM-optimized memtable** that extends the write buffer via remote memory and adds **shard-level flush offloading** to idle compute/memory nodes, increasing flush parallelism and accelerating L0 compaction.
- Implemented **cache-enhanced read delegation**, a local key–offset cache with RDMA-assisted remote access, delivering up to **5.2×** faster point/range queries; outperformed Disaggregated-RocksDB, CaaS-LSM, and Nova-LSM in benchmarks.

### Optimizing Approximate Nearest Neighbor Search (ANNS) in Vector Databases with Disaggregated Memory

- Built vector indexes in **Disaggregated Memory (DM)** and served lookups via **RDMA**, decoupling compute from memory to provide memory-class access at cluster scale.
- Designed an **adaptive RDMA read-granularity** controller that tunes block size/prefetch by ANN phase (candidate selection, refinement, re-ranking), reducing over-fetch and bus contention.
- Overlapped ANN compute with **RDMA I/O** via prefetch, request coalescing, and double-buffered queues, saturating the NIC while keeping CPUs busy, yielding higher throughput and lower tail latency.

### LLM-Agent Driven Auto-Tuning Framework for Heterogeneous Data Systems

- Built an autonomous agentic LLM framework with **tool use, planning, reasoning, and decision-making** to tune configurations for **RocksDB, CacheLib, and MySQL InnoDB** end-to-end (workload benchmarking, log analysis, configuration optimization).
- Designed a **modular multi-agent** pipeline (Executor, Extractor, Reflector, Searcher) combining **insight-driven tree search** and **feedback-validation** for safe configuration tuning, mitigating hallucinations and unsafe actions.
- Achieved up to **5.75×** throughput and **88%** lower latency on real-world benchmarks, exceeded state-of-the-art baselines by **63%** and prior LLM methods by **38%**.

### ACADEMIC SERVICES

---

- **Program Committee Member (Demo Track)**, ACM SIGSpatial 2025
- **Artifact Evaluation Committee Member**, FAST 2025; SIGMOD 2025
- **Journal Reviewer**, ACM Transactions on Software Engineering and Methodology (TOSEM); IEEE Transactions on Information Forensics and Security (TIFS)

### GRANTS & AWARDS

---

- **ASU Fulton Fellowship**, Arizona State University (2023–2027)
- **Academic Exchange Funding**, Southeast University (top 5% selection; 46/1,100+ students), 2019
- **National First Prize**, *China Undergraduate Mathematical Contest in Modeling (CUMCM)* (Top 1%; 292/28,665 teams)

### SKILLS

---

- **Programming Languages:** C/C++, Python, Java, SQL, MATLAB, Solidity
- **Frameworks & Tools:** Linux, Docker, Kafka, Flink, RocksDB, Velox, PyTorch, NumPy, pandas, OpenCV, STL, Git, AWS, Google Cloud Platform (GCP), Spring Boot, Spring Cloud, Ethereum, Confluence, Jama, Jira
- **Domains:** Distributed Systems, Key–Value Stores, Storage Systems, High-Performance Networking (RDMA, CXL), Database Systems, Blockchain, Machine Learning, Data Science, Software Engineering